
Pseudolikelihood inference for Gibbsian T-tessellations... and point processes

Kiên KIÊU and Katarzyna ADAMCZYK-CHAUVAUD

Rapport technique 2015-2, novembre 2015

UR1404 Mathématiques et Informatique Appliquées du Génome à l'Environnement
INRA

Domaine de Vilvert

78352 Jouy-en-Josas Cedex

France

<http://maiage.jouy.inra.fr>

© 2015 INRA



Kiên KIÊU

UR1404 MaIAGE, INRA, domaine de Vilvert, 78352 Jouy-en-Josas Cedex
kien.kieu@jouy.inra.fr

Katarzyna ADAMCZYK-CHAUVAT

UR1404 MaIAGE, INRA, domaine de Vilvert, 78352 Jouy-en-Josas Cedex

Abstract

Recently a new class of planar tessellations, named T-tessellations, was introduced. Splits, merges and a third local modification named flip were shown to be sufficient for exploring the space of T-tessellations. Based on these local transformations and by analogy with point process theory, tools Campbell measures and a general simulation algorithm of Metropolis-Hastings-Green type were translated for random T-tessellations.

The current report is concerned with parametric inference for Gibbs models of T-tessellations. The estimation criterion referred to as the pseudolikelihood is derived from Campbell measures of random T-tessellations and the Kullback-Leibler divergence. A detailed algorithm for approximating the pseudolikelihood maximum is provided. A simulation study seems to show that bias and variability of the pseudolikelihood maximum decrease when the tessellated domain grows in size.

In the last part of the report, it is shown that an analogous approach based on the Campbell measure and the KL divergence when applied to point processes leads to the well-known pseudo-likelihood introduced by Besag. More surprisingly, the binomial regression method recently proposed by Baddeley and his co-authors for computing the pseudolikelihood maximum can be derived using the same approach starting from a slight modification of the Campbell measure.

Keywords: Stochastic geometry, tessellation, point process, pseudolikelihood, Campbell measure, Kullback-Leibler divergence.

Contents

1	Introduction	5
2	Pseudolikelihood for T-tessellations	5
2.1	Gibbsian T-tessellations	5
2.2	Pseudolikelihood definition and properties	8
2.3	The canonical exponential family case	12
2.4	Algorithm for approximating the pseudolikelihood maximum	15
2.5	Simulation study	17
2.6	Fitting a wrong model	21
3	Pseudolikelihood for point processes	24
3.1	Finite point processes	26
3.2	Contrast function and pseudolikelihood	26
3.3	A weaker but more tractable contrast function	27
4	Conclusion	30
A	Extension of the Kullback-Leibler divergence to non probability measures	32
B	Subconfigurations of a Poisson point process	33
C	A random T-tessellation is determined by its Papangelou conditional intensities	34

1 Introduction

Recently, a new class of planar tessellations was introduced [12]: Gibbsian T-tessellations. Briefly, a T-tessellation is a tessellation with only vertices at the intersection of three edges, two of them being aligned. In [12], it was shown that the space of T-tessellations can be explored using three simple and local operators: splits, merges and flips. A model was proposed to be considered as a completely random T-tessellation model (CRTT). And Gibbsian variants were discussed. Also tools and results such as Campbell measures, Papangelou kernels, Georgii-Nguyen-Zessin formulae were translated from point process theory to T-tessellations. A Metropolis-Hastings-Green algorithm based on splits, merges and flips was derived for the simulation of Gibbsian T-tessellations. Such an algorithm defines a Split, Merge and Flip (SMF) Markov chain converging to the Gibbs distribution to be simulated. As a companion tool, a C++ library named LiTe [1] was developed. Using LiTe, it is possible to simulate a rather large class of T-tessellation Gibbs models.

The main purpose of this report is to introduce a method for estimating the parameters of a T-tessellation Gibbs model. This inference method is based on a so-called pseudolikelihood. The approach developed here involves Campbell measures and the Kullback-Leibler divergence. It is shown that the well-known pseudolikelihood introduced by Besag [5] for point process can also be built from the Campbell measure and the Kullback-Leibler divergence.

The proposal for the pseudolikelihood for Gibbsian T-tessellations is provided in Section 2.2. Section 2.3 focuses on the special case of a canonical exponential family. Practical computation of the pseudolikelihood maximum is discussed in Section 2.4. The behaviour of the maximum pseudolikelihood estimator is investigated based on some numerical simulations in Section 2.5. Section 2.6 is devoted to an unexpected drawback of the pseudolikelihood.

In Section 3.2, it is shown that a similar approach based on the Campbell measure and the Kullback-Leibler divergence provides a new way to obtain the widely used pseudolikelihood for point processes. Even the recently logistic regression method for computing the pseudolikelihood estimator can be obtained in such a way as shown in Section 3.3.

Appendix A is a technical prerequisite where the Kullback-Leibler divergence is extended to non-probability measures. Appendix B provides a result about the mean of a functional summed over all subpatterns of a finite Poisson point process. Such a result is used for proving that the distribution of a Gibbsian T-tessellation is fully characterized by its Papangelou conditional intensities (detailed proof provided in Appendix C).

2 Pseudolikelihood for T-tessellations

2.1 Gibbsian T-tessellations

A general theory of planar T-tessellations was sketched in [12]. At the time being, the theory is limited to bounded domains. Let $D \subset \mathbb{R}^2$ be a convex bounded polygon.

Below, D is referred to as the domain. The space of T -tessellations of D is denoted by \mathcal{T} . A T -tessellation has only convex cells and can also be seen as a collection of segments. A segment is a maximal union of aligned and connected edges. There are two classes of segments: non-blocking (single edge) and blocking (multiple edges) segments. Three local operators allow to explore \mathcal{T} :

Split a cell is divided by a new segment.

Merge a non-blocking segment is removed.

Flip an edge at the end of a blocking segment is removed while the incident blocked segment is extended.

Note that the inverse of a split is a merge (and vice versa), while the inverse of a flip is a flip. Below \mathbb{S}_T , \mathbb{M}_T and \mathbb{F}_T refer to the spaces of splits, merges and flips that can be applied to tessellation T . Note that \mathbb{S}_T is continuous while \mathbb{M}_T and \mathbb{F}_T are finite. In [12], uniform measures with infinitesimal elements denoted by ds , dm and df are defined. The last two measures are just counting measures. The measure ds has a total mass equal to the sum of the perimeters of T cells, denoted by $u(T)$, divided by π . The probability measure $\pi ds/u(T)$ on \mathbb{S}_T is characterized by the two following facts:

- Cells of T are split with probabilities proportional to their perimeters.
- Given a cell C of T is split, the splitting line is a uniform and isotropic random line hitting C .

The so-called CRTT model introduced in [12] plays the role of the Poisson point process in point process theory. Its distribution is denoted by μ .

Definition 1. For a given non-negative real number ν , the CRTT on D with parameter ν is defined by the following distribution on \mathcal{T} :

$$\mu_\nu(dT) \propto I\{T \in \mathcal{T}(L)\} \frac{\nu^n}{n!} dl_1 \dots dl_n, \quad \forall L = \{l_1, \dots, l_n\}, \quad (1)$$

where $\mathcal{T}(L)$ is the subset of T -tessellations with segments supported by the line pattern L .

For sake of brevity, we write μ the distribution of the CRTT with parameter $\nu = 1$. CRTT stands for Completely Random T -Tessellation.

All other models considered below have a density with respect to μ . Their distributions can be written as

$$P(dT) = h(T) \mu(dT). \quad (2)$$

In general, h is known analytically only up to a multiplicative normalizing constant. Even for the CRTT where h is constant, the normalizing constant is not known at the moment. This is why inference by maximum likelihood is not straightforward.

The density h is said to be hereditary in the sense defined in [12] if:

- For every pair (T, s) , $T \in \mathcal{T}$, $s \in \mathbb{S}_T$, $h(T) = 0 \Rightarrow h(sT) = 0$.
- For every pair (T, f) , $T \in \mathcal{T}$, $f \in \mathbb{F}_T$, $h(T) = 0 \Rightarrow h(fT) = 0$.

The condition was erroneously stated in [12] with inversed implications.

In [12], two reduced Campbell measures for random T -tessellations were introduced. The split Campbell measure of a random T -tessellation \mathbf{T} maps any function $\phi : \mathcal{C}_s = \{(s, T) : T \in \mathcal{T}, s \in \mathbb{S}_T\} \rightarrow \mathbb{R}$ to

$$C_s^!(\phi) = \mathbb{E} \sum_{m \in \mathbb{M}_{\mathbf{T}}} \phi(m^{-1}, m\mathbf{T}). \quad (3)$$

The split Campbell measure is very similar to the Campbell measure of a point process. Indeed splitting and merging are analog to insertion and removal of points.

The flip Campbell measure is more specific to T -tessellations but its definition is formally very close to the definition of the split Campbell measure. The flip Campbell measure of \mathbf{T} maps any function $\phi : \mathcal{C}_f = \{(f, T) : T \in \mathcal{T}, f \in \mathbb{F}_T\} \rightarrow \mathbb{R}$ to

$$C_f^!(\phi) = \mathbb{E} \sum_{f \in \mathbb{F}_{\mathbf{T}}} \phi(f^{-1}, f\mathbf{T}). \quad (4)$$

It was shown in [12] that if $\mathbf{T} \sim h\mu$ where h is hereditary then

$$C_s^!(ds, dT) = \lambda_s(s, T) ds P(dT), \quad (5)$$

$$C_f^!(f, dT) = \lambda_f(f, T) df P(dT), \quad (6)$$

where

$$\lambda_s(s, T) = \frac{h(sT)}{h(T)}, \quad (7)$$

$$\lambda_f(f, T) = \frac{h(fT)}{h(T)}, \quad (8)$$

for P -almost all $T \in \mathcal{T}$, almost all $s \in \mathbb{S}_T$ and all $f \in \mathbb{F}_T$. The densities λ_s and λ_f of the reduced Campbell measures are called the split and flip Papangelou conditional intensities of \mathbf{T} . Here we will need a third one: the merge Papangelou conditional intensity defined by

$$\lambda_m(m, T) = \frac{h(mT)}{h(T)}. \quad (9)$$

Note that

$$\lambda_m(m, T) = \lambda_s(m^{-1}, mT)^{-1}. \quad (10)$$

The split and flip Papangelou conditional intensities uniquely determine the density h when it is hereditary. A proof of the result below is given in Appendix C.

Proposition 1. *Let $\mathbf{T} \sim P = h\mu$ where h is supposed to be hereditary. The density h is uniquely determined up to a μ -null set by any pair (λ_s, λ_f) of split and flip Papangelou conditional intensities of \mathbf{T} .*

As this report focuses on parametric inference, most often a parametric family of distributions P_θ on \mathcal{T} will be considered. The issue addressed in this report is to estimate θ given a realization of P_θ .

Since a general simulation algorithm of Metropolis-Hastings-Green is provided in [12], inference by Monte-Carlo Maximum Likelihood is possible, see e.g. [9]. But its practical implementation can be quite demanding. So easier and more automatized inference methods are of interest especially in contexts where estimation has to be repeatedly applied with different models or on large data sets.

The pseudolikelihood approach developped below is based on Campbell measures and Kullback-Leibler divergence. Using discretization as in [5] for point processes is hardly concevable for tessellations.

2.2 Pseudolikelihood definition and properties

Consider a parametric family of distributions P_θ on \mathcal{T} , $\theta \in \Theta$. Every P_θ is supposed to be absolutely continuous with respect to μ and the densities of the P_θ 's are denoted by h_θ . The model is supposed to be identifiable:

$$\theta = \theta^* \Leftrightarrow h_\theta(T) = h_{\theta^*}(T) \text{ for } \mu\text{-almost all } T \in \mathcal{T}. \quad (11)$$

The densities h_θ are supposed to be hereditary. The split, merge and flip Papangelou intensities are denoted $\lambda_{s,\theta}$, $\lambda_{m,\theta}$ and $\lambda_{f,\theta}$.

We are now ready for the pseudolikelihood definition.

Definition 2. *Given an observed T -tessellation T , the log-pseudolikelihood is defined as*

$$\begin{aligned} \text{LPL}(\theta|T) = & - \sum_{m \in \mathbb{M}_T} \log \lambda_{m,\theta}(m, T) - \int_{\mathbb{S}_T} \lambda_{s,\theta}(s, T) \, ds - \\ & \sum_{f \in \mathbb{F}_T} \log \lambda_{f,\theta}(f, T) - \sum_{f \in \mathbb{F}_T} \lambda_{f,\theta}(f, T). \end{aligned} \quad (12)$$

The maximum pseudolikelihood estimator is defined as

$$\hat{\theta} = \arg \max_{\theta \in \Theta} \text{LPL}(\theta|T). \quad (13)$$

Similarly to the log-likelihood, the log-pseudolikelihood can be derived as an empirical estimator of a contrast function. As a starting point, instead of P_θ , consider the split and flip Campbell measures:

$$C_{s,\theta}^!(ds, dT) = \lambda_{s,\theta}(s, T) \, ds P_\theta(dT), \quad (14)$$

$$C_{f,\theta}^!(f, dT) = \lambda_{f,\theta}(f, T) \, df P_\theta(dT), \quad (15)$$

Compared to P_θ , the Campbell measures have densities that do not involve unknown normalizing constants. It should be noticed however that the dominating measure above

depends on θ . This is why we introduce further measures closely related to the Campbell measures. For every pair (θ, θ^*) , consider the measures $\alpha_{s,\theta,\theta^*}$ on \mathcal{C}_s and $\alpha_{f,\theta,\theta^*}$ on \mathcal{C}_f defined by

$$\alpha_{s,\theta,\theta^*}(ds, dT) = \lambda_{s,\theta}(s, T) ds P_{\theta^*}(dT), \quad (16)$$

$$\alpha_{f,\theta,\theta^*}(df, dT) = \lambda_{f,\theta}(f, T) df P_{\theta^*}(dT). \quad (17)$$

The measures $\alpha_{s,\theta,\theta^*}$ and $\alpha_{f,\theta,\theta^*}$ can be viewed as candidate Campbell measures. When $\theta = \theta^*$, they are indeed Campbell measures.

For the time being, let us focus on the case where $\alpha_{s,\theta,\theta^*} \gg \alpha_{s,\theta^*,\theta^*}$ and $\alpha_{f,\theta,\theta^*} \gg \alpha_{f,\theta^*,\theta^*}$, then

$$\frac{d\alpha_{s,\theta^*,\theta^*}}{d\alpha_{s,\theta,\theta^*}}(s, T) = \frac{\lambda_{s,\theta^*}(s, T)}{\lambda_{s,\theta}(s, T)}, \quad \frac{d\alpha_{f,\theta^*,\theta^*}}{d\alpha_{f,\theta,\theta^*}}(f, T) = \frac{\lambda_{f,\theta^*}(f, T)}{\lambda_{f,\theta}(f, T)}.$$

It is standard in statistics to quantify the dissemblance between distributions using the Kullback-Leibler divergence. Since the measures $\alpha_{s,\theta,\theta^*}$ and $\alpha_{f,\theta,\theta^*}$ are not probability measures, using the Kullback-Leibler divergence requires some adaptation. Appendix A provides such an extension to non-negative non-probability measures. The extended Kullback-Leibler divergence of $\alpha_{s,\theta,\theta^*}$ from $\alpha_{s,\theta^*,\theta^*}$ is defined as

$$\begin{aligned} \text{KL}(\alpha_{s,\theta^*,\theta^*}, \alpha_{s,\theta,\theta^*}) \\ = \int_{\mathcal{C}_s} \left(\frac{\lambda_{s,\theta^*}(s, T)}{\lambda_{s,\theta}(s, T)} \log \frac{\lambda_{s,\theta^*}(s, T)}{\lambda_{s,\theta}(s, T)} + 1 - \frac{\lambda_{s,\theta^*}(s, T)}{\lambda_{s,\theta}(s, T)} \right) \alpha_{s,\theta,\theta^*}(ds, dT) \end{aligned}$$

Note that for every $\theta \in \Theta$

$$\text{KL}(\alpha_{s,\theta^*,\theta^*}, \alpha_{s,\theta,\theta^*}) \geq 0$$

and that the extended Kullback-Leibler divergence cancels for $\theta = \theta^*$.

Replacing in the right-hand side $\alpha_{s,\theta,\theta^*}$ by its definition (16), one gets

$$\begin{aligned} \text{KL}(\alpha_{s,\theta^*,\theta^*}, \alpha_{s,\theta,\theta^*}) \\ = \int_{\mathcal{C}_s} \left(\lambda_{s,\theta^*}(s, T) \log \frac{\lambda_{s,\theta^*}(s, T)}{\lambda_{s,\theta}(s, T)} + \lambda_{s,\theta}(s, T) - \lambda_{s,\theta^*}(s, T) \right) ds P_{\theta^*}(dT). \end{aligned}$$

The first part of the integral can be rewritten using the expression of the split Campbell measure (14):

$$\begin{aligned} \text{KL}(\alpha_{s,\theta^*,\theta^*}, \alpha_{s,\theta,\theta^*}) \\ = \int_{\mathcal{C}_s} \log \frac{\lambda_{s,\theta^*}(s, T)}{\lambda_{s,\theta}(s, T)} C_{s,\theta^*}^!(ds, dT) + \int_{\mathcal{C}_s} (\lambda_{s,\theta}(s, T) - \lambda_{s,\theta^*}(s, T)) ds P_{\theta^*}(dT). \end{aligned}$$

Define

$$M_s(\theta, \theta^*) = \int_{\mathcal{C}_s} \log \lambda_{s,\theta}(s, T) C_{s,\theta^*}^!(ds, dT) - \int_{\mathcal{C}_s} \lambda_{s,\theta}(s, T) ds P_{\theta^*}(dT).$$

The extended Kullback-Leibler divergence of $\alpha_{s,\theta,\theta^*}$ from $\alpha_{s,\theta^*,\theta^*}$ can be written as

$$\text{KL}(\alpha_{s,\theta,\theta^*}, \alpha_{s,\theta,\theta^*}) = M_s(\theta^*, \theta^*) - M_s(\theta, \theta^*).$$

Only the term $M_s(\theta, \theta^*)$ is informative about the divergence of $\alpha_{s,\theta,\theta^*}$ from $\alpha_{s,\theta^*,\theta^*}$. Note that M_s does not have a constant sign. Instead, for any given θ^* , $M_s(\cdot, \theta^*)$ has a global maximum at θ^* .

Now suppose that θ^* is unknown and that a realization T of P_{θ^*} is available instead. In view of definition (3), the first integral of M_s can be estimated by

$$\sum_{m \in \mathbb{M}_T} \log \lambda_{s,\theta}(m^{-1}, mT).$$

From the equivalence (10), one can rewrite the estimator as

$$- \sum_{m \in \mathbb{M}_T} \log \lambda_{m,\theta}(m, T).$$

Deriving an estimator of the second integral of M_s is straightforward. The mean under P_{θ^*} of the integral

$$\int_{\mathbb{S}_T} \lambda_{s,\theta}(s, T) ds$$

is equal to the second integral involved in the expression of $M_s(\theta, \theta^*)$ above.

Similarly considering the extended KL divergence of $\alpha_{f,\theta,\theta^*}$ from $\alpha_{f,\theta^*,\theta^*}$ leads to the function:

$$M_f(\theta, \theta^*) = \int_{\mathcal{T}} \int_{\mathbb{F}_T} \log \lambda_{f,\theta}(f, T) C_{f,\theta^*}^!(df, dT) - \int_{\mathcal{T}} \int_{\mathbb{F}_T} \lambda_{f,\theta}(f, T) df P_{\theta^*}(dT).$$

Furthermore, $M_f(\theta, \theta^*)$ can be estimated without bias by

$$- \sum_{f \in \mathbb{F}_T} \log \lambda_{f,\theta}(f, T) - \sum_{f \in \mathbb{F}_T} \lambda_{f,\theta}(f, T).$$

We can now state a first general result about the mean log-pseudolikelihood

Proposition 2. *If $\mathbf{T} \sim P_{\theta^*}$, then for every $\theta \in \Theta$*

$$\begin{aligned} E_{\theta^*} \text{LPL}(\theta | \mathbf{T}) &= M_s(\theta, \theta^*) + M_f(\theta, \theta^*), \end{aligned} \tag{18}$$

$$= M_s(\theta^*, \theta^*) + M_f(\theta^*, \theta^*) - \text{KL}(\alpha_{s,\theta^*,\theta^*}, \alpha_{s,\theta,\theta^*}) - \text{KL}(\alpha_{f,\theta^*,\theta^*}, \alpha_{f,\theta,\theta^*}). \tag{19}$$

Proof. The Proposition has already been proved under the assumption that $\alpha_{s,\theta,\theta^*} \gg \alpha_{s,\theta^*,\theta^*}$ and $\alpha_{f,\theta,\theta^*} \gg \alpha_{f,\theta^*,\theta^*}$. Let us consider the case where either $\alpha_{s,\theta,\theta^*}$ does not dominate $\alpha_{s,\theta^*,\theta^*}$ or $\alpha_{f,\theta,\theta^*}$ does not dominate $\alpha_{f,\theta^*,\theta^*}$. For instance, consider the first case. There exists a non-null subset, with respect to the measure $ds P_{\theta^*}(dT)$, of \mathcal{C}_s where $\lambda_{s,\theta}$ cancels while λ_{s,θ^*} does not. In view of expression (5), this is equivalent with

the existence of a non-null subset, with respect to the measure $C_{s,\theta^*}^!$, where $\lambda_{s,\theta}$ cancels. Hence if $\alpha_{s,\theta,\theta^*}$ does not dominate $\alpha_{s,\theta^*,\theta^*}$, then $M_s(\theta, \theta^*)$ is equal to $-\infty$. Since, by definition, $\text{KL}(\alpha_{s,\theta^*,\theta^*}, \alpha_{s,\theta,\theta^*}) = \infty$, the equality between (18) and (19) holds. A similar rationale applies for flips instead of splits. \square

A straightforward consequence of Proposition 2 is the following.

Corollary 1. *If $\mathbf{T} \sim P_{\theta^*}$, then for every $\theta \in \Theta$,*

$$E_{\theta^*} \text{LPL}(\theta|\mathbf{T}) \leq E_{\theta^*} \text{LPL}(\theta^*|\mathbf{T})$$

Proof. In view of equations (18–19), we have for every $\theta \in \Theta$

$$E_{\theta^*} \text{LPL}(\theta|\mathbf{T}) = E_{\theta^*} \text{LPL}(\theta^*|\mathbf{T}) - \text{KL}(\alpha_{s,\theta^*,\theta^*}, \alpha_{s,\theta,\theta^*}) - \text{KL}(\alpha_{f,\theta^*,\theta^*}, \alpha_{f,\theta,\theta^*}). \quad (20)$$

Since the extended Kullback-Leibler divergences are non-negative, the inequality follows. \square

The following result shows that $-(M_s(., \theta^*) + M_f(., \theta^*))$ is a contrast function in the sense that its global minimum at θ^* is strict.

Proposition 3. *If $\mathbf{T} \sim P_{\theta^*}$, then θ^* is a strict global maximum of the mean log-pseudolikelihood:*

$$\forall \theta \in \Theta \setminus \{\theta^*\}, \quad E_{\theta^*} \text{LPL}(\theta|\mathbf{T}) < E_{\theta^*} \text{LPL}(\theta^*|\mathbf{T}).$$

Proof. It remains to show that

$$E_{\theta^*} \text{LPL}(\theta|\mathbf{T}) = E_{\theta^*} \text{LPL}(\theta^*|\mathbf{T}) \Rightarrow \theta = \theta^*.$$

Consider a θ such that the mean log-pseudolikelihoods are equal. First, observe that in view of Equation (20), both Kullback-Leibler divergences must cancel. This implies that $\alpha_{s,\theta,\theta^*}$ (respectively $\alpha_{f,\theta,\theta^*}$) must dominate $\alpha_{s,\theta^*,\theta^*}$ (respectively $\alpha_{f,\theta^*,\theta^*}$), otherwise one of the KL-divergences is infinite. Hence the mean log-pseudolikelihoods can be equal only if $\alpha_{s,\theta,\theta^*} \gg \alpha_{s,\theta^*,\theta^*}$ and $\alpha_{f,\theta,\theta^*} \gg \alpha_{f,\theta^*,\theta^*}$. From now on, we focus on that case.

The Kullback-Leibler divergence of $\alpha_{s,\theta,\theta^*}$ from $\alpha_{s,\theta^*,\theta^*}$ cancels if and only if $\alpha_{s,\theta,\theta^*} = \alpha_{s,\theta^*,\theta^*}$ except on a $\alpha_{s,\theta,\theta^*}$ -null subset $E \subset \mathcal{C}_s$. In view of definition (16), E is $\alpha_{s,\theta,\theta^*}$ -null if and only there exists a pair (E_1, E_2) of subsets of \mathcal{C}_s such that $E = E_1 \cup E_2$, E_1 is null with respect to the measure $dsP_{\theta^*}(dT)$ and $\lambda_{s,\theta}$ cancels on E_2 . The domination $\alpha_{s,\theta,\theta^*} \gg \alpha_{s,\theta^*,\theta^*}$ implies that λ_{s,θ^*} must cancel too on E_2 . Therefore $\alpha_{s,\theta,\theta^*}$ and $\alpha_{s,\theta^*,\theta^*}$ coincide on E_2 . More generally, both measures equal except on E_1 which is null with respect to $dsP_{\theta^*}(dT)$. In turn, this implies that $\lambda_{s,\theta}$ and λ_{s,θ^*} equal for P_{θ^*} -almost all $T \in \mathcal{T}$ and almost all splits $s \in \mathbb{S}_T$. Hence, λ_{s,θ^*} is a split Papangelou conditional intensity for P_{θ^*} .

Similarly, it can be shown that $\lambda_{f,\theta}$ is a flip Papangelou conditional intensity for P_{θ^*} . From Proposition 1, it follows that h_θ is a density of P_{θ^*} and that it equals h_{θ^*} μ -almost everywhere. Since the model is supposed to be identifiable, θ must equal θ^* . \square

2.3 The canonical exponential family case

Consider the case of a canonical exponential family:

$$P_\theta(dT) \propto \exp(\theta^T t(T)) \mu(dT), \quad (21)$$

wher $t(T) \in \mathbb{R}^d$ is a vector of statistics of tessellation T . Then we have

$$\begin{aligned} \lambda_{s,\theta}(s, T) &= \exp(\theta^T t(s, T)), & t(s, T) &= t(sT) - t(T), \\ \lambda_{m,\theta}(m, T) &= \exp(\theta^T t(m, T)), & t(m, T) &= t(mT) - t(T), \\ \lambda_{f,\theta}(f, T) &= \exp(\theta^T t(f, T)), & t(f, T) &= t(fT) - t(T). \end{aligned}$$

The log-pseudolikelihood writes

$$\begin{aligned} \text{LPL}(\theta|T) &= -\theta^T \sum_{m \in \mathbb{M}_T} t(m, T) - \int_{\mathbb{S}_T} \exp(\theta^T t(s, T)) ds - \\ &\quad \theta^T \sum_{f \in \mathbb{F}_T} t(f, T) - \sum_{f \in \mathbb{F}_T} \exp(\theta^T t(f, T)). \end{aligned} \quad (22)$$

The gradient of the log-pseudolikelihood has the following form

$$\begin{aligned} \nabla \text{LPL}(\theta|T) &= - \sum_{m \in \mathbb{M}_T} t(m, T) - \int_{\mathbb{S}_T} t(s, T) \exp(\theta^T t(s, T)) ds - \\ &\quad \sum_{f \in \mathbb{F}_T} t(f, T) - \sum_{f \in \mathbb{F}_T} t(f, T) \exp(\theta^T t(f, T)). \end{aligned} \quad (23)$$

The Hessian matrix of LPL has elements of the form

$$\begin{aligned} H_{ij}(\text{LPL})(\theta|T) &= - \int_{\mathbb{S}_T} t_i(s, T) t_j(s, T) \exp(\theta^T t(s, T)) ds - \\ &\quad \sum_{f \in \mathbb{F}_T} t_i(f, T) t_j(f, T) \exp(\theta^T t(f, T)). \end{aligned}$$

That is

$$\begin{aligned} H(\text{LPL})(\theta|T) &= - \int_{\mathbb{S}_T} t(s, T) t(s, T)^T \exp(\theta^T t(s, T)) ds - \\ &\quad \sum_{f \in \mathbb{F}_T} t(f, T) t(f, T)^T \exp(\theta^T t(f, T)) \end{aligned}$$

As a consequence the log-pseudolikelihood is concave. Strict concavity holds if for every $y \in \mathbb{R}^d$

$$\begin{aligned} y^T H(\text{LPL})(\theta|T) y &= - \int_{\mathbb{S}_T} (y^T t(s, T))^2 \exp(\theta^T t(s, T)) ds - \\ &\quad \sum_{f \in \mathbb{F}_T} (y^T t(f, T))^2 \exp(\theta^T t(f, T)) < 0. \end{aligned}$$

Define

$$\mathbb{S}_T^* = \{s \in \mathbb{S}_T : t(s, T) > -\infty\}, \quad \mathbb{F}_T^* = \{f \in \mathbb{F}_T : t(f, T) > -\infty\}. \quad (24)$$

In the quadratic form above, the integral can be restricted to \mathbb{S}_T^* and the sum to \mathbb{F}_T^* . The log-pseudolikelihood is strictly concave if at least one of the two following conditions is fulfilled.

Condition 1. *The subset $t(\mathbb{F}_T^*, T)$ of \mathbb{R}^d is not included in a vector hyperplane.*

Condition 2. *For any subset B of \mathbb{S}_T^* almost equal to \mathbb{S}_T^* , $t(B, T)$ is not included in a vector hyperplane of \mathbb{R}^d .*

Above B is almost equal to \mathbb{S}_T^* if and only if

$$\int_{\mathbb{S}_T} I_{\mathbb{S}_T^* \setminus B}(s) ds = 0.$$

Condition 1 ensures that for every $y \in \mathbb{R}^d$, there exists at least a flip $f \in \mathbb{F}_T^*$ such $t(f, T)$ is not orthogonal to y . Similarly Condition 2 ensures that there exists a non-null subset of \mathbb{S}_T^* where $t(s, T)$ is not orthogonal to y .

Example 1. Consider the distribution defined by the density

$$h_\theta(T) \propto \exp\left(\overset{\circ}{n}_s(T)\theta\right) \quad (25)$$

where $\overset{\circ}{n}_s(T)$ refers to the number of internal segments of T . Above the density h_θ is defined up to a unknown normalizing constant depending on θ . This distribution defines a CRTT with intensity controlled by parameter θ . In [12], the CRTT model is introduced together with another parametrization based on $\tau = \exp(\theta)$. The parameter τ can be considered as a scaling parameter: increasing τ is equivalent to increasing the tessellated domain D . Let us come back to parametrization (25) which defines a canonical exponential family with

$$t(T) = \overset{\circ}{n}_s(T),$$

and thus

$$\begin{aligned} t(s, T) &= 1, \\ t(m, T) &= -1, \\ t(f, T) &= 0. \end{aligned}$$

The log-pseudolikelihood writes

$$\text{LPL}(\theta|T) = |\mathbb{M}_T| \theta - \frac{u(T)}{\pi} \exp(\theta) - |\mathbb{F}_T|.$$

The log-pseudolikelihood can be rewritten in terms of the number of internal segments of T :

$$\text{LPL}(\theta|T) = \overset{\circ}{n}_{\text{s,nb}}(T) \theta - \frac{u(T)}{\pi} \exp(\theta) - 2\overset{\circ}{n}_{\text{s,b}}(T),$$

where $\overset{\circ}{n}_{\text{s,nb}}(T)$ is the number of non-blocking internal segments of T and $\overset{\circ}{n}_{\text{s,b}}(T)$ the number of blocking internal segments. The function above is strictly concave (negative second derivative). It has therefore a unique maximum. In order to identify the maximum, one solves

$$\text{LPL}'(\hat{\theta}|T) = \overset{\circ}{n}_{\text{s,nb}}(T) - \frac{u(T)}{\pi} \exp(\hat{\theta}) = 0.$$

This yields

$$\hat{\theta} = \log \frac{\overset{\circ}{n}_{\text{s,nb}}(T) \pi}{u(T)}.$$

That is

$$\hat{\tau} = \frac{\overset{\circ}{n}_{\text{s,nb}}(T) \pi}{u(T)}.$$

Note that for the CRTT model, claimed to be the analog of the Poisson point process, the pseudolikelihood maximum estimator is not the same as the likelihood estimator. Indeed for the CRTT model the likelihood has the form

$$\frac{\tau^{\overset{\circ}{n}_{\text{s}}(T)}}{c(\tau)},$$

where $c(\tau)$ is a normalization constant (unknown so far). It follows that the likelihood maximum estimator of τ is a function (still unknown) of the number $\overset{\circ}{n}_{\text{s}}(\tau)$ of internal segments of the tessellation T .

The estimators derived from the likelihood and the pseudolikelihood do not coincide for completely random T -tessellations contrary to what happens for Poisson point processes.

Example 2. Consider the distribution defined (up to an unknown normalizing constant) by the density

$$h_{\theta}(T) \propto \exp \left(\overset{\circ}{n}_{\text{s}}(T) \theta_1 - a^2(T) \theta_2 \right), \quad (26)$$

where $a^2(T)$ is the sum of T cell squared areas. Given the number of cells the statistics $a^2(T)$ is large when the cells have heterogeneous areas. Large values of θ_2 favour tessellations with homogeneous cells (with respect to their areas). Again we have a canonical exponential family with

$$t(T) = \begin{pmatrix} \overset{\circ}{n}_{\text{s}}(T) \\ -a^2(T) \end{pmatrix}.$$

Expressions for $t(s, T)$, $t(m, T)$ and $t(f, T)$ can be worked out. But this does not lead to a simple analytical expression of the pseudolikelihood.

2.4 Algorithm for approximating the pseudolikelihood maximum

A simple approach for approximating the log-pseudolikelihood consists in discretizing the integral on \mathbb{S}_T involved in the pseudolikelihood. Below we consider the framework of a canonical exponential family. The log-pseudolikelihood and its gradient have the forms given by Equations (22) and (23).

Let S be a finite set of so-called *dummy* splits of T . The number of elements of S is denoted $|S|$. The log-pseudolikelihood (22) is replaced by

$$\begin{aligned} \text{LPL}_d(\theta|T; S) = & -\theta^T \sum_{m \in \mathbb{M}_T} t(m, T) - \frac{u(T)}{\pi|S|} \sum_{s \in S} \exp(\theta^T t(s, T)) - \\ & \theta^T \sum_{f \in \mathbb{F}_T} t(f, T) - \sum_{f \in \mathbb{F}_T} \exp(\theta^T t(f, T)). \end{aligned} \quad (27)$$

The approximation concerns only the integral over \mathbb{S}_T involved in (22) which is replaced by a weighted discrete sum. All other terms involving merges and flips are exact. Notice that if S is a sample of splits drawn uniformly from \mathbb{S}_T , then $\text{LPL}_d(\theta|T; S)$ is an unbiased estimator of $\text{LPL}(\theta|T)$.

Remark 1. *In the special case of the CRTT model, the approximation is exact. Since $t(s, T) = 1$ for any split s , the sum over S in (27) writes*

$$\sum_{s \in S} \exp(\theta^T t(s, T)) = |S| \exp(\theta).$$

Thus the term involving dummy splits becomes

$$\frac{u(T)}{\pi} \exp(\theta).$$

Compare with Example 1.

A straightforward approximation of the log-pseudolikelihood gradient is

$$\begin{aligned} \nabla \text{LPL}_d(\theta|T; S) = & - \sum_{m \in \mathbb{M}_T} t(m, T) - \frac{u(T)}{\pi|S|} \sum_{s \in S} t(s, T) \exp(\theta^T t(s, T)) - \\ & \sum_{f \in \mathbb{F}_T} t(f, T) - \sum_{f \in \mathbb{F}_T} t(f, T) \exp(\theta^T t(f, T)). \end{aligned} \quad (28)$$

Similarly, the log-pseudolikelihood Hessian can be approximated by

$$\begin{aligned} H(\text{LPL}_d)(\theta|T; S) = & -\frac{u(T)}{\pi|S|} \sum_S t(s, T) t(s, T)^T \exp(\theta^T t(s, T)) - \\ & \sum_{f \in \mathbb{F}_T} t(f, T) t(f, T)^T \exp(\theta^T t(f, T)) \end{aligned} \quad (29)$$

Algorithm 1 Algorithm NOIS (Newton optimization and increasing splitting) for approximating the pseudolikelihood maximum.

Require: a T-tessellation T , the stepsize ϵ to be used in θ 's update, the tolerance parameter δ

- 1: $m \leftarrow$ number of non-blocking segments of T
 - 2: $\theta \leftarrow$ null vector
 - 3: Compute $\sum t(m, T)$ for $m \in \mathbb{M}_T$
 - 4: Compute $t(f, T)$ for all $f \in \mathbb{F}_T$
 - 5: Compute the sum of the $t(f, T)$'s
 - 6: $S \leftarrow$ sample of m independent and uniform splits s of T
 - 7: Compute the $t(s, T)$'s
 - 8: **repeat**
 - 9: $G \leftarrow \nabla \text{LPL}_d(\theta|T; S)$ as given by (28)
 - 10: $H \leftarrow H(\text{LPL}_d)(\theta|T; S)$ as given by (29)
 - 11: $\theta \leftarrow \theta + \epsilon H^{-1}G$
 - 12: Add m independent and uniform splits of T to S
 - 13: Update the list of $t(s, T)$'s where $s \in S$
 - 14: $L \leftarrow \text{LPL}_d(\theta|T; S)$ as given by (27)
 - 15: $\Delta \leftarrow$ variation of L
 - 16: **until** $|\Delta| \leq \delta (|L| + \delta)$.
-

The integral on splits in the log-pseudolikelihood (22) can be written as an expectation with respect to the uniform distribution on the space \mathbb{S}_T of splits. The estimation criterion consists of deterministic terms and of an expectation with respect to a random variable that can be simulated. Stochastic approximation provides a number of algorithms for optimizing such type of criterion. However the present framework is rather simple here. The expectation is taken under a fixed distribution not depending on the searched parameter.

From stochastic approximation we keep the alternation between the update of the parameter and the estimation of the expectation. An iteration of the algorithm consists of the two following stages:

- Update of θ according to Newton's optimization method based on the log-pseudolikelihood Hessian approximation.
- Increment of the sample S of dummy splits.

The full procedure is detailed in Algorithm 1. It is referred to as NOIS (Newton optimization and increasing splitting) algorithm below. Some remarks:

- The memory space used by the algorithm for storing the $t(s, T)$'s grows along iterations.
- The time for computing

$$\sum_{s \in S} \exp(\theta^T t(s, T))$$

grows along iterations too.

- One could add one dummy split to S at each iteration, but the convergence of the sum on dummy splits is slow and the trajectories of the θ_n 's are then rather irregular.
- The stopping criterion is inspired by the R function `optim`.
- When the parametric family includes the CRTT distributions, one may start with the maximum pseudolikelihood estimate under the CRTT model. Without loss of generality, one may assume that the first component of $t(T)$ is equal to $\overset{\circ}{n}_s(T)$. The initial value of θ writes:

$$\left(\log \left(\overset{\circ}{n}_{s,nb}(T) \pi \right) / u(T), 0, \dots, 0 \right)^T.$$

Example 3. Consider the model with the following non-normalized density

$$h_\theta(T) \propto \exp \left(\theta_1 \overset{\circ}{n}_s(T) - \theta_2 \sum_{v \in V(T)} \left(\frac{\pi}{2} - \phi(v) \right) \right), \quad (30)$$

where $0 \leq \phi(v) \leq \pi/2$ denotes the acute angle between the two segments meeting at vertex v . When $\theta_2 > 0$, T -tessellations with almost rectangular cells are favoured by the model. A realization of such a model with $\theta_1 = 2.49$ and $\theta_2 = 2.5$ is shown in Figure 1 (top). Figure 1 (bottom) shows an example of parameter trajectory obtained with the NOIS algorithm. Some details:

- The observed tessellation was generated by a SMF Markov chain starting from the empty tessellation and after a burnin stage of 1,000 iterations.
- A sample of 1,000 dummy splits was used for computing the discrete approximation of the pseudolikelihood and of its gradient.
- As a starting point for the optimization, the pseudolikelihood estimator assuming the CRTT model was used.
- Ten iterations have been performed before the stopping criterion was met with $\delta = 0.01$. Since the T -tessellation used for inference had 31 non-blocking segments, the NOIS algorithm only used $10 \times 31 = 310$ dummy splits.

2.5 Simulation study

At the time being, it is not possible to investigate the asymptotic behaviour of the maximum pseudolikelihood estimator. Gibbs models of T -tessellations are defined only in bounded domains. Extensions to the whole plane, existence and unicity of Gibbs measures are still open questions. In order to provide some hints on the behaviour of




Figure 1: Top: T-tessellation generated by the model with penalization on acute angles between segments. Bottom: maximum pseudolikelihood estimation. Background: discrete approximation of the log-pseudolikelihood (increasing from red to white). Grey arrows: approximated log-pseudolikelihood gradient. Circles: intermediate estimates found by NOIS algorithm. Cross: true parameters.

the maximum pseudolikelihood estimator, we proceeded to a simulation study. Three models of random T-tessellations and three domain sizes for each model are considered.

The first model is the CRTT model defined in Example 1. The second model is defined in Example 3 and favour T-tessellations with perpendicular segments. The last model favours T-tessellation with homogeneous cell areas and is defined in Example 2. The three models are referred to as CRTT, angle and area models below. Parameter values are given in Table 1. Each model has been simulated on three squared domains with varying sizes. The domain sizes were chosen in order to emphasize their effect on the estimator precision. Domain side lengths are provided in Table 1.

Table 1: Simulation settings. For each model, θ values are given. Tessellated domains are squares with side lengths w_1 , w_2 and w_3 . Numbers of iterations for the burnin stage and between sampled tessellations (sampling period) are also provided.

model	θ_1	θ_2	w_1	w_2	w_3	burnin	period
CRTT	0.64		2.5	1.75	1	12500	3704
angle	2.49	2.5	2.5	1.75	1.4	30000	9473
area	0.53	835.2	2.8	1.9	1.5	11000	7223

The maximum pseudolikelihood estimator has been computed on samples of 500 T-tessellations for each model and domain size. In order to generated the T-tessellation, SMF Markov chains has been used as described in [12]. As for any Metropolis-Hastings algorithm, burnin and sampling period (number of Markov chain iterations between consecutive sampled items) have to be determined.

The burnin duration was tuned empirically by monitoring the evolution of the model energy. It was chosen as the number of iterations required for energy stabilization (visual assessment).

Concerning the sampling period, an automatic procedure was used. Along an SMF chain curse, segments are born (splits), change their length (flips) or die (merge). Thus a segment has a survival time expressed in terms of number of iterations. As a criterion, the renewal rate of segments was considered. The sampling period was chosen long enough so that 75 % of segments living at the beginning of the period had died at the end of the period. Sampling periods were approximated using the RLiTe function `getSMFSamplingPeriod`.

Note that burnin duration and sampling periods were tuned for each model on the largest domain to be considered. Simulations on smaller domains were performed using the SMF Markov chains with parameters chosen for the largest domains.

First sample times determined from the burnin duration and the sampling period are shown for the three models simulated in their largest domains in Figure 2.

In order to compute estimates using the LiTe class `PLInferenceNOIS`, one needs to specify the tolerance and the maximal number of iterations. Values are given in Table 2. Estimation for the CRTT model is apart since the maximum pseudolikelihood has




Figure 2: Simulation of the CRTT (top row), angle (middle row) and area (bottom row) models. Left: first tessellation sampled at the end of the burnin stage. Right: energy evolution during the first SMF iterations. Vertical lines: iterations where tessellations are sampled.

Table 2: Estimation settings used by the NOIS algorithm. For each model, the values of the tolerance parameter δ and the maximal number of NOIS updates.

model	tolerance	max. iterations
CRTT		0
angle	0.005	150
area	-0.005	100

an explicit simple form. It can be computed using method `Run` of class `PLInferenceNOIS` with a null maximum number of iterations. The maximal number of iterations for the angle model was chosen large enough so that the stopping criterion could be met in most cases. Concerning the area model, it turned out that the stopping criterion was met within very few iterations (sometimes only one for the smallest model). This resulted in rather biased estimates for the parameter θ_1 . This poor performance may be due to flat pseudolikelihood functions as illustrated below. In order to bypass this problem, the tolerance parameter δ is chosen negative so that the stopping criterion cannot be met and the maximum number of NOIS iterations is chosen large.

Estimate distributions for the CRTT model are shown in Figure 3 (top row) as boxplots. Dispersion is moderate and decreases when the domain size increases. Small bias: it is not visible as soon as the total number of cells is larger than a hundred. For fifty cells (smallest domain size), a small negative bias cannot be excluded.

Estimate distributions for the angle model are shown in Figure 3 (middle row). As for the CRTT model, the estimate dispersion and bias seem to decrease when the domain size increases. The estimate distributions show some asymmetry.

Let us consider the area model. As mentioned above, the stopping criterion of the NOIS algorithm does not seem to work for the area model. This problem is illustrated based on a simulated tessellation. The pseudolikelihood is approximated using 5,000 dummy splits, see Figure 4. From now on this approximation is considered as exact. In the region around the maximum, the log-pseudolikelihood is rather flat: relative variations of the log-pseudolikelihood are small. As a consequence, the NOIS algorithm stops too early and find an estimate a little bit too far from the true maximum. A possible enhancement of the stopping criterion: stop when the pseudolikelihood and the estimates do not vary enough any more. An alternative solution could be to reparametrize the density (26) in order to get two statistics of T with more comparable values.

Estimation for the area model behaves like the two previous models, see Figure 3 (bottom row). Both dispersion and bias decrease when the domain size increases.

2.6 Fitting a wrong model

Let us consider an observed T-tessellation approximating agricultural plots located around Selommès village (Loir-et-Cher department, France). This tessellation is shown in Figure 5 (center). The tessellated domain has a width of about 5.3 km. There are 211

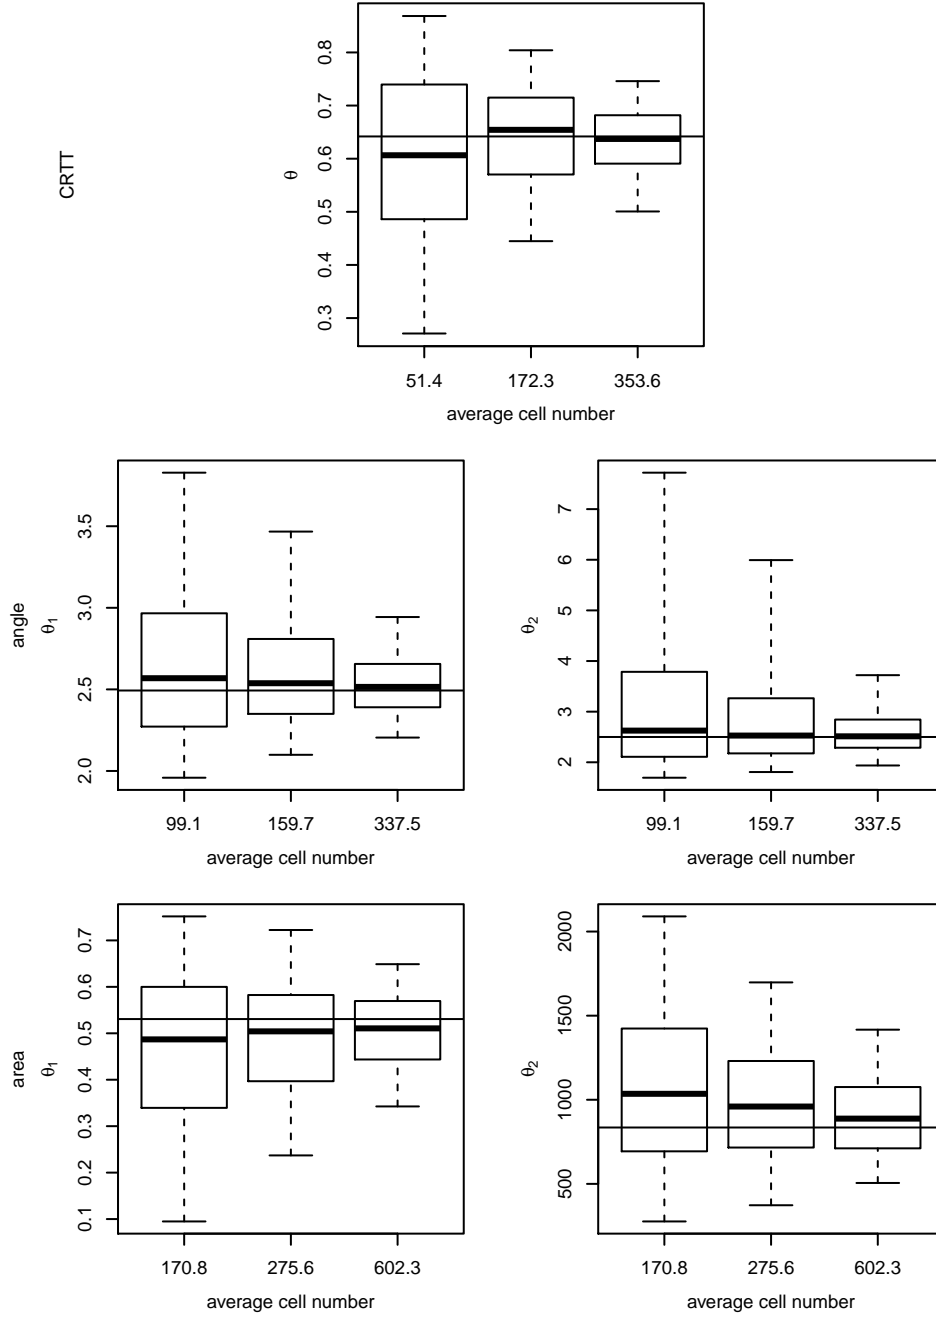


Figure 3: Maximum pseudolikelihood estimates (500 replicates) for the three models (CRTT, angle and area). Abscissa: mean number of cells for each of the three domain size. For each model and domain size, the estimate distribution is represented as a boxplot. First and last deciles are shown as the whisker ends. True parameter values represented as horizontal lines.




Figure 4: Pseudolikelihood of a simulated tessellation under the area model. Intermediate estimates are shown as plus signs when found before the stopping criterion was fulfilled. Further estimates are shown as crosses. The location of the true maximum is shown as a circle.

agricultural plots (cells). Choosing a good model for representing such a tessellation is an open research problem far beyond the scope of this paper. As a toy model, consider the CRTT model. It is not plausible but rather simple. Our observed T-tessellation has 86 internal non-blocking segments. The sum of all cell perimeters is equal to 213.2 km. Thus the estimate of θ is equal to $\hat{\theta} = 0.24$.

Simulations of T-tessellations under the CRTT model with parameter $\theta = 0.24$ are shown in Figure 5. The most striking feature of simulated tessellations is their heterogeneity compared to the observed T-tessellation. This is expected since the CRTT model yields completely random T-tessellation while the observed tessellation shows some kind of regularity. But there is another important difference between the observed and the simulated tessellations: there are more cells in the latter.

This difference is unexpected. Such a behaviour does not occur with point processes. In the point process framework, the situation is as follows. Consider a given point pattern of say n points. Suppose that the homogeneous Poisson model is chosen for fitting. The intensity parameter is estimated by n divided by the area of the considered domain. This estimate is obtained both as the maximum pseudolikelihood and as the maximum likelihood estimators. Therefore simulating the fitted Poisson model would yield point patterns of size n in expectation. Somehow pseudolikelihood fitting leads to the best Poisson point process representing the observed point pattern. In the T-tessellation framework, the fitted model seems to be biased for reasons that remain to be investigated.

3 Pseudolikelihood for point processes

Pseudolikelihood was first introduced by Besag for spatial random fields with discrete support [5], later for a Strauss point process [6]. The extension to a rather general class of models for point processes was provided by Jensen and Møller [11]. General convergence results (consistency, central limit theorem) were obtained by Billiot et al. [7]. In this report, we focus on point processes within a bounded domain denoted by D . The way the pseudolikelihood is introduced here is original. Previously the point process was discretized by countings on a mesh (a square lattice for example). This yields a random field with discrete support. By letting the tile size decrease to zero, one obtains the point process pseudolikelihood as a limit of the random field pseudolikelihood (the limit does not depend on the considered mesh). Similarly to random T-tessellations, one can derive the pseudolikelihood starting from the Kullback-Leibler divergence and candidate Campbell measures. Such an approach is developed below. It should be noticed that inference based on pseudolikelihood can also be considered as a particular instance of large class of methods so-called Takacs-Fiksel methods [8].

Below only the main steps of the pseudolikelihood derivations are sketched. Full details can be recovered by analogy with the pseudolikelihood derivation for random T-tessellations.

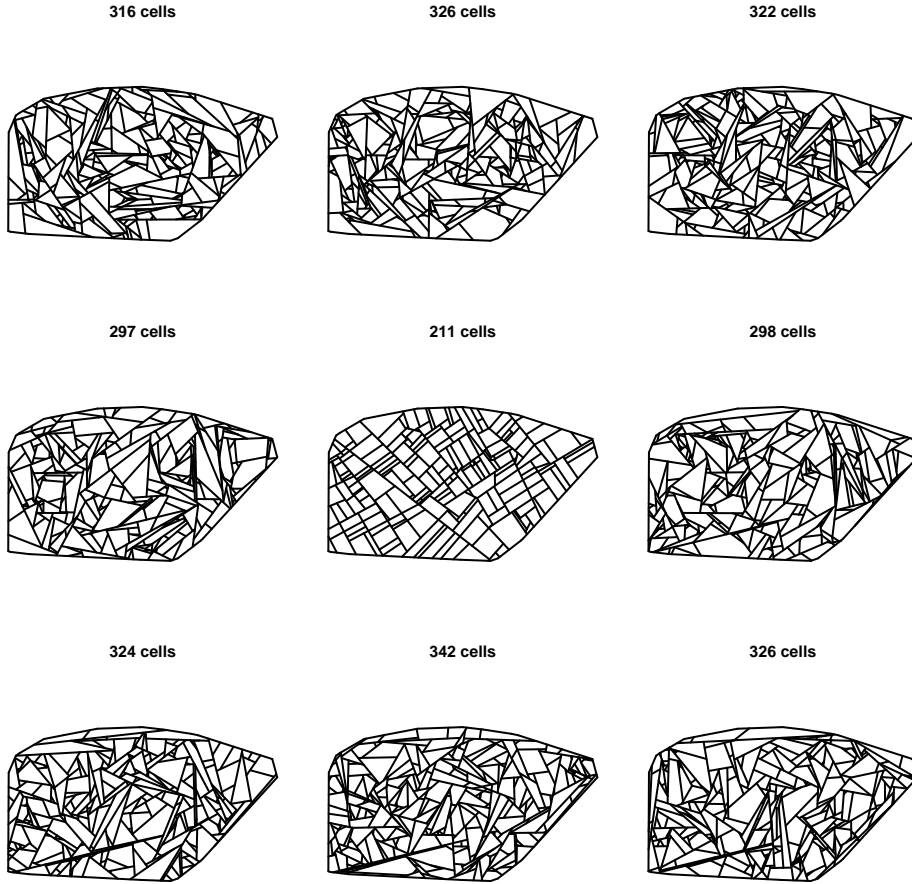


Figure 5: Center: a T-tessellation supposed to be a realization of the CRTT model with unknown parameter. Around: realizations of the CRTT model fitted using the maximum pseudolikelihood estimator.

3.1 Finite point processes

Let \mathcal{D} denote the space of finite point patterns in a planar compact domain D . Considering only bounded domains simplifies the presentation and makes the analogy with random T-tessellations more straightforward.

Any point process \mathbf{X} defines a reduced Campbell measure $C^!$ on $D \times \mathcal{D}$ which maps any measurable function $\phi : D \times \mathcal{D} \rightarrow \mathbb{R}$ to

$$C^!(\phi) = \mathbb{E} \sum_{x \in \mathbf{X}} \phi(x, \mathbf{X} \setminus \{x\}). \quad (31)$$

The Campbell measure above is very similar to the split Campbell measure defined for random T-tessellations. A split of a T-tessellation is the insertion of a non-blocking segment. Its analog for point patterns is the insertion of a new point. A merge of a T-tessellation is the removal of a non-blocking segment. Its analog for point patterns is the removal of an (arbitrary) point of the pattern.

Let μ denote the distribution of a Poisson point process with intensity 1 on D . Consider a distribution $P \ll \mu$ on \mathcal{D} with density h :

$$P(dX) = h(X) \mu(dX).$$

Let us suppose that h is hereditary :

$$\forall x \in D, \forall X \in \mathcal{D}, \quad h(X) = 0 \Rightarrow h(X \cup \{x\}) = 0. \quad (32)$$

Then the reduced Campbell measure is absolutely continuous with respect to the product of the Lebesgue measure on \mathbb{R}^2 with P and its density denoted λ with respect to the product measure is defined as the Papangelou conditional intensity given by

$$C^!(dx, dX) = \lambda(x; X) dx P(dX)$$

The Papangelou conditional intensity has the following expression

$$\lambda(x; X) = \frac{h(X \cup \{x\})}{h(X)}. \quad (33)$$

Note that one can find in the literature expressions of the Papangelou conditional intensity slightly different from (33). For instance, for $x \in X$, Baddeley and Turner [3] define $\lambda(x; X)$ as $h(X)/h(X \setminus x)$ instead of 1. The different versions of the Papangelou conditional intensity are consistent with (33) as long as they coincide with (33) almost-everywhere with respect to the product of the Lebesgue measure on \mathbb{R}^2 with P .

3.2 Contrast function and pseudolikelihood

Consider a parametric family of distributions P_θ on \mathcal{D} . It is supposed that every $P_\theta \ll \mu$ and has a hereditary density h_θ with respect to μ . Also the model is supposed to be

identifiable. The associated Papangelou conditional intensities are denoted λ_θ . For every pair (θ, θ^*) , consider the measure $\alpha_{\theta, \theta^*}$ on $D \times \mathcal{D}$ defined by

$$\alpha_{\theta, \theta^*}(dx, dX) = \lambda_\theta(x, X) dx P_{\theta^*}(dX). \quad (34)$$

The Kullback-Leibler divergence of $\alpha_{\theta, \theta^*}$ from $\alpha_{\theta^*, \theta^*}$ can be expressed as

$$\text{KL}(\alpha_{\theta^*, \theta^*}, \alpha_{\theta, \theta^*}) = M(\theta^*, \theta^*) - M(\theta, \theta^*),$$

where

$$M(\theta, \theta^*) = \int_{D \times \mathcal{D}} \log \lambda_\theta(x, X) C_{\theta^*}^1(dx, dX) - \int_{D \times \mathcal{D}} \lambda_\theta(x, X) dx P_{\theta^*}(dX) \quad (35)$$

Note that $M(\cdot, \theta^*)$ has a global maximum at $\theta = \theta^*$. Similarly to random T -tessellations, it can be shown that the maximum is strict.

Let X be a realization of $\mathbf{X} \sim P_{\theta^*}$, then

$$\text{LPL}(\theta|X) = \sum_{x \in X} \log \lambda_\theta(x, X \setminus \{x\}) - \int_D \lambda_\theta(x, X) dx \quad (36)$$

is an unbiased estimator of $M(\theta, \theta^*)$. Equation (36) coincides with the log-pseudolikelihood introduced by Besag [5, 6, 11] for point processes. The maximum pseudolikelihood estimator is

$$\hat{\theta} = \arg \max_{\theta} \text{LPL}(\theta|X).$$

Generally no analytical expression of the integral involved in the log-pseudolikelihood (36) is available. Thus computing $\hat{\theta}$ is not straightforward. The most commonly used method for computing $\hat{\theta}$ is the Poisson regression method. This method was first proposed by Berman and Turner [4] for the inference of heterogeneous Poisson point processes. The extension to a larger class of Gibbs processes was developed by Baddeley and Turner [3].

3.3 A weaker but more tractable contrast function

Below we introduce a family of contrast functions. They can be considered as approximations of the contrast function $-M$ associated with the pseudolikelihood. Their main feature is that their minimum is easily computed using standard tools of statistics.

Let ρ be an integrable function on D . In practice, ρ may be taken proportional to the indicator function of D . Define

$$\alpha_{\theta, \theta^*}^\rho(dx, dX) = (\lambda_\theta(x, X) + \rho(x)) dx P_{\theta^*}(dX). \quad (37)$$

When $\rho \equiv 0$, the measure defined in (37) is just $\alpha_{\theta, \theta^*}$ as defined in Equation (34). Consider the function

$$\text{KL}(\alpha_{\theta^*, \theta^*}, \alpha_{\theta, \theta^*}) - \text{KL}(\alpha_{\theta^*, \theta^*}^\rho, \alpha_{\theta, \theta^*}^\rho). \quad (38)$$

This function cancels at $\theta = \theta^*$ and it is non-negative according to Lemma 1. The following inequalities hold:

$$0 \leq \text{KL}(\alpha_{\theta^*, \theta^*}, \alpha_{\theta, \theta^*}) - \text{KL}(\alpha_{\theta^*, \theta^*}^\rho, \alpha_{\theta, \theta^*}^\rho) \leq \text{KL}(\alpha_{\theta^*, \theta^*}, \alpha_{\theta, \theta^*}).$$

Heuristically, the function (38) tends to the upper bound when ρ tends to infinity while it tends to the lower bound when ρ tends to zero. Hence the function (38) of θ

- Approximates $\text{KL}(\alpha_{\theta^*, \theta^*}, \alpha_{\theta, \theta^*})$ as ρ tends to infinity.
- Is less discriminative than $\text{KL}(\alpha_{\theta^*, \theta^*}, \alpha_{\theta, \theta^*})$ for a finite ρ .

Simple computations show that up to some additive terms not depending on θ the KL-divergence can be written as

$$\begin{aligned} \text{KL}(\alpha_{\theta^*, \theta^*}^\rho, \alpha_{\theta, \theta^*}^\rho) &= - \int_{D \times \mathcal{D}} \log(\lambda_\theta(u, X) + \rho(u)) C_{\theta^*}^!(du, dX) \\ &\quad - \int_D \int_{\mathcal{D}} \rho(u) \log(\lambda_\theta(u, X) + \rho(u)) P_{\theta^*}(dX) du + \int_D \int_{\mathcal{D}} \lambda_\theta(u, X) P_{\theta^*}(dX) du \\ &\quad + \dots \end{aligned}$$

Thus the difference (38) can be written as

$$\begin{aligned} &- \int_{D \times \mathcal{D}} \log \frac{\lambda_\theta(u, X)}{\lambda_\theta(u, X) + \rho(u)} C_{\theta^*}^!(du, dX) + \int_D \int_{\mathcal{D}} \rho(u) \log(\lambda_\theta(u, X) + \rho(u)) P_{\theta^*}(dX) du \\ &\quad + \dots \end{aligned}$$

As a consequence the M function introduced in (35) can be redefined as follows

$$\begin{aligned} M^\rho(\theta, \theta^*) &= \int_{D \times \mathcal{D}} \log \frac{\lambda_\theta(u, X)}{(\lambda_\theta(u, X) + \rho(u))} C_{\theta^*}^!(du, dX) \\ &\quad - \int_D \int_{\mathcal{D}} \rho(u) \log(\lambda_\theta(u, X) + \rho(u)) P_{\theta^*}(dX) du. \quad (39) \end{aligned}$$

The function M^ρ has a global maximum $\theta = \theta^*$. Below it is rewritten as an expectation in order to derive an unbiased estimator. The first term can be rewritten as

$$\int_{D \times \mathcal{D}} \log \frac{\lambda_\theta(u, X)}{(\lambda_\theta(u, X) + \rho(u))} C_{\theta^*}^!(du, dX) = \mathbb{E}_{\theta^*} \sum_{x \in \mathbf{X}} \log \frac{\lambda_\theta(x, \mathbf{X} \setminus \{x\})}{\lambda_\theta(x, \mathbf{X} \setminus \{x\}) + \rho(x)},$$

where \mathbb{E}_{θ^*} is the mean when $\mathbf{X} \sim P_{\theta^*}$. In order to rewrite the second term, consider a point process \mathbf{Y} on D with intensity ρ . Then we have

$$\int_D \int_{\mathcal{D}} \rho(u) \log(\lambda_\theta(u, X) + \rho(u)) P_{\theta^*}(dX) du = \mathbb{E} \mathbb{E}_{\theta^*} \sum_{y \in \mathbf{Y}} \log(\lambda_\theta(y, \mathbf{X}) + \rho(y)), \quad (40)$$

where the first mean is taken with respect to \mathbf{Y} . Therefore an unbiased estimator of M^ρ is

$$\sum_{x \in \mathbf{X}} \log \frac{\lambda_\theta(x, \mathbf{X} \setminus \{x\})}{\lambda_\theta(x, \mathbf{X} \setminus \{x\}) + \rho(x)} - \sum_{y \in \mathbf{Y}} \log (\lambda_\theta(y, \mathbf{X}) + \rho(y)). \quad (41)$$

The unbiasedness of (41) can be checked as follows:

- The mean with respect to $\mathbf{X} \sim P_{\theta^*}$ of the first sum is equal to the first integral in (39) by definition of the reduced Campbell measure (31).
- The mean with respect to both \mathbf{Y} and $\mathbf{X} \sim P_{\theta^*}$ of the second sum is equal to the second integral in (39) in view of identity (40).

Let $\hat{\theta}$ be the estimator of θ^* that maximizes the above criterion with respect to θ . It maximizes also

$$\text{LPL}^\rho(\theta|X) = \sum_{x \in X} \log \frac{\lambda_\theta(x, X \setminus \{x\})}{\lambda_\theta(x, X \setminus \{x\}) + \rho(x)} + \sum_{y \in Y} \log \frac{\rho(y)}{\lambda_\theta(y, X) + \rho(y)} \quad (42)$$

since both criteria differ only by a term not depending on θ . In order to obtain an even more homogeneous expression, assume that the conditional probability given \mathbf{X} that $\mathbf{X} \cap \mathbf{Y} = \emptyset$ is zero. Then almost surely $\mathbf{X} = \mathbf{X} \setminus \{y\}$ and the estimating criterion can be rewritten as

$$\text{LPL}^\rho(\theta|X) = \sum_{x \in X} \log \frac{\lambda_\theta(x, X \setminus \{x\})}{\lambda_\theta(x, X \setminus \{x\}) + \rho(x)} + \sum_{y \in Y} \log \frac{\rho(y)}{\lambda_\theta(y, X \setminus \{y\}) + \rho(y)}. \quad (43)$$

The estimation criterion above was first proposed by Baddeley et al. [2]. Our derivation here is different: based on a contrast function rather than a score function.

Thus we get the log-likelihood of a binomial regression. The sample to be considered is $t_z, z \in X \cup Y$ with $t_z = 1$ when $z \in X$. And the success probabilities are of the form

$$\frac{\lambda_\theta(z, X \setminus \{z\})}{\lambda_\theta(z, X \setminus \{z\}) + \rho(z)}.$$

Let us focus on the case of an exponential family. Then we have

$$\begin{aligned} \Pr[t_z = 1] &= \frac{\exp(\theta^T t(z, X \setminus \{z\}))}{\exp(\theta^T t(z, X \setminus \{z\})) + \rho(z)}, \\ &= \frac{1}{1 + \exp(\log \rho(z) - \theta^T t(z, X \setminus \{z\}))}. \end{aligned}$$

That is

$$\text{logit } \Pr[t_z = 1] = \log \rho(z) - \theta^T t(z, X \setminus \{z\}).$$

Therefore in order to compute $\hat{\theta}$ one can use computing tools already available for the binomial regression with a logistic link function.

4 Conclusion

In our paper [12], a new class of stochastic models for planar tessellations was introduced. Theoretical concepts and results mainly inspired by point process theory were provided. Also a Metropolis-Hastings-Green type algorithm was proposed for simulating Gibbs models. The current report is a first step towards parametric inference from observed data.

The key ingredient for introducing the pseudolikelihood for Gibbsian T-tessellations is the Kullback-Leibler divergence and contrast function theory. Applying such an approach to the density h would lead to the classical likelihood which is unpracticable because of the unknown normalizing function involved in h . The trick here is to consider Campbell measures instead. As shown in this report, applying a similar approach to point process theory yields the well-known pseudolikelihood introduced by Besag [5] who discretized point processes into spatial random fields with discrete supports. It was also shown that the approximation of the maximum pseudolikelihood based on a binomial regression [2] can be derived from a contrast function involving the Kullback-Leibler divergence and the Campbell measure.

Whether such an approach can be extended to other types of random spatial structures is a open question. In our opinion, the key point is to identify basic operators that allow to explore the whole space of spatial structures. Based on such operators, one may define one or several Campbell measures leading in turn to a pseudolikelihood.

Alternative inference methods are also of interest. As a variant, one could consider two contrast functions associated with splits and merges on one side and on flips on the other side. This would lead to two pseudolikelihoods and two estimators. Another possibility is Monte-Carlo Maximum Likelihood (MCML), see e.g. [10]. Implementing MCML is possible using the simulation algorithm proposed in [12]. Bayesian approaches would deserve some investigations too.

In this report, the behaviour of the maximum pseudolikelihood estimator was investigated by means of simulations showing that bias and dispersion decrease as the number of cells in the observed tessellation increases. Theoretical studies would be of interest too, in particular asymptotics. However such studies require the development of a theoretical framework where random T-tessellation would be extended to the whole plane (currently random T-tessellations are bounded to a compact polygon).

The intriguing behaviour of the maximum pseudolikelihood when fitting a wrong model to observed data should be investigated either theoretically or empirically.

Applying the maximum pseudolikelihood to real data may require some preprocessing. As an example, consider the maximum pseudolikelihood for the CRTT model. Its computation requires the determination of the number of non-blocking segments and of the total cell perimeter. In practice, a T-tessellation is often provided as a list of polygons (cells). Therefore one has to compute the number of non-blocking segments from the list of polygons. This is not a trivial task both because real tessellations may not be exact T-tessellations and because of unavoidable numerical errors related to number representations in digital computing. The library LiTe provides some methods which

may help for transforming sets of polygons into T-tessellations.

References

- [1] K. Adamczyk-Chauvat and K. Kiêu. LiTe. <http://kien-kieu.github.io/lite>, 2015.
- [2] A. Baddeley, J.-F. Coeurjolly, E. Rubak, and R. Waagepetersen. Logistic regression for spatial Gibbs point processes. *Biometrika*, 101(2):377–392, 2014.
- [3] A. Baddeley and R. Turner. Practical maximum likelihood for spatial point patterns. *Australian & New Zealand Journal of Statistics*, 42(3):283–322, 2000.
- [4] M. Berman and T. R. Turner. Approximating point process likelihoods with GLIM. *Journal of the Royal Statistical Society, Series C, Applied Statistics*, 41(1):31–38, 1992.
- [5] J. Besag. Statistical analysis of non-lattice data. *The Statistician*, 24(3):179–195, 1975.
- [6] J. Besag. Some methods of statistical analysis for spatial data. *Bulletin of the International Statistical Institute*, 47:77–91, 1977.
- [7] J.-M. Billiot, J.-F. Coeurjolly, and R. Drouilhet. Maximum pseudolikelihood estimator for exponential family models of marked Gibbs point processes. *Electronic Journal of Statistics*, 2:234–264, 2008.
- [8] J.-F. Coeurjolly, D. Dereudre, R. Drouilhet, and F. Lavancier. Takacs-Fiksel method for stationary marked Gibbs point processes. *Scandinavian Journal of Statistics*, 39(3):416–443, 2011.
- [9] C. Geyer. Likelihood inference for spatial point processes. In O. Barndorff-Nielsen, W. Kendall, and M. van Lieshout, editors, *Stochastic Geometry - Likelihood and Computation*, volume 80 of *Monographs on Statistics and Applied Probability*, pages 79–140. Chapman and Hall/CRC, 1999.
- [10] C. J. Geyer. Markov chain Monte Carlo maximum likelihood. In *Computing Science and Statistics: Proceedins of the 23rd Symposium on the Interface*, pages 156–163, 1991.
- [11] J. L. Jensen and J. Møller. Pseudolikelihood for exponential family models of spatial point processes. *The Annals of Applied Probability*, 3:445–461, 1991.
- [12] K. Kiêu, K. Adamczyk-Chauvat, H. Monod, and R. S. Stoica. A completely random T-tessellation model and Gibbsian extensions. *Spatial Statistics*, 6:118–138, 2013.

A Extension of the Kullback-Leibler divergence to non probability measures

Let α and β be two finite non-negative measures on a measurable space \mathcal{E} . Let the extended Kullback-Leibler divergence of β from α be defined as

$$\text{KL}(\alpha, \beta) = \begin{cases} \int_{\mathcal{E}} \left(\frac{d\alpha}{d\beta}(z) \log \frac{d\alpha}{d\beta}(z) + 1 - \frac{d\alpha}{d\beta}(z) \right) \beta(dz), & \text{if } \beta \gg \alpha, \\ +\infty & \text{otherwise.} \end{cases} \quad (44)$$

In the equation above, $0 \log 0$ is assumed to be 0. If α and β have the same total mass, Equation (44) simplifies into

$$\text{KL}(\alpha, \beta) = \int_{\mathcal{E}} \frac{d\alpha}{d\beta}(z) \log \frac{d\alpha}{d\beta}(z) \beta(dz).$$

The simplification above applies especially when α and β are probability measures. The classical Kullback-Leibler divergence is a particular case of the extended one.

Since the function

$$x \rightarrow x \log x + 1 - x$$

is non-negative, the extended Kullback-Leibler divergence defined in Equation (44) is also non-negative. Furthermore the function above cancels only at $x = 1$. It follows that the extended Kullback-Leibler divergence cancels only when α equals β except on a β -null subset of \mathcal{E} .

A small lemma that will be used further in this report.

Lemma 1. *Let α , β and γ be non-negative finite measures on a measurable space \mathcal{E} . Then the following inequality holds*

$$\text{KL}(\alpha, \beta) \geq \text{KL}(\alpha + \gamma, \beta + \gamma).$$

The result is rather intuitive: $\alpha + \gamma$ and $\beta + \gamma$ are more similar than α and β .

Proof. First consider the case where β does not dominate α . Then $\text{KL}(\alpha, \beta) = +\infty$ and the inequality holds.

Now assume that $\beta \gg \alpha$. Let E be a measurable subset $E \subset \mathcal{E}$ such that $(\beta + \gamma)(E) = 0$. Both β and γ cancel on E . Since $\beta \gg \alpha$, α cancels on E too. It follows that $(\alpha + \gamma)(E) = 0$. Hence $(\beta + \gamma)(E) = 0$ implies that $(\alpha + \gamma)(E) = 0$, i.e. $\beta + \gamma \gg \alpha + \gamma$.

Let μ be a measure dominating α , β and γ . Let a , b and c be the densities of α , β and γ with respect to μ . We have

$$\begin{aligned} & \text{KL}(\alpha + \gamma, \beta + \gamma) \\ &= \int_{\mathcal{E}} \left(\frac{a(z) + c(z)}{b(z) + c(z)} \log \frac{a(z) + c(z)}{b(z) + c(z)} + 1 - \frac{a(z) + c(z)}{b(z) + c(z)} \right) (b(z) + c(z)) \mu(dz), \\ &= \int_{\mathcal{E}} \left((a(z) + c(z)) \log \frac{a(z) + c(z)}{b(z) + c(z)} + b(z) - a(z) \right) \mu(dz). \end{aligned}$$

The difference between divergences can be written as

$$\begin{aligned} & \text{KL}(\alpha, \beta) - \text{KL}(\alpha + \gamma, \beta + \gamma) \\ &= \int_{\mathcal{E}} \left(a(z) \log \frac{a(z)}{b(z)} - (a(z) + c(z)) \log \frac{a(z) + c(z)}{b(z) + c(z)} \right) d\mu(z). \end{aligned}$$

We have to check that the integrand is non-negative. The non-negativity results from the following inequality

$$x \log \frac{x}{y} \geq (x + h) \log \frac{x + h}{y + h}, \quad \text{for all } h \geq 0.$$

The latter inequality is a consequence of the decreasing of the function

$$\binom{x}{y} \rightarrow x \log \frac{x}{y}$$

when moving away from the diagonal $y = x$. The decreasing is proved by

$$\left\langle \nabla x \log \frac{x}{y}, \begin{pmatrix} 1 \\ 1 \end{pmatrix} \right\rangle = \left\langle \begin{pmatrix} \log \frac{x}{y} + 1 \\ -\frac{x}{y} \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \end{pmatrix} \right\rangle = 1 - \frac{x}{y} + \log \frac{x}{y} \leq 0.$$

Above $\langle \cdot, \cdot \rangle$ denotes the scalar product. □

B Subconfigurations of a Poisson point process

Lemma 2. *Let \mathbf{X} be a Poisson point process with intensity measure ν on a space D . The total mass $\nu(D)$ is supposed to be finite. Then, for any measurable real-valued function ϕ defined on the space of finite point patterns in D , we have*

$$\mathbb{E} \sum_{Y \subset \mathbf{X}} \phi(Y) = \exp(\nu(D)) \mathbb{E} \phi(\mathbf{X}). \quad (45)$$

Proof. Using the general expression of the distribution of a Poisson point process, one gets

$$\mathbb{E} \sum_{Y \in \mathbf{X}} \phi(Y) = e^{-\nu(D)} \sum_{n \geq 0} \frac{1}{n!} \int_{D^n} \sum_{Y \subset \{x_1, \dots, x_n\}} \phi(Y) \nu^n(dx_1 \dots dx_n).$$

A subset of a point pattern can be mapped into a subset of indices:

$$\mathbb{E} \sum_{Y \in \mathbf{X}} \phi(Y) = e^{-\nu(D)} \sum_{n \geq 0} \frac{1}{n!} \sum_{I \subset \{1, \dots, n\}} \int_{D^n} \phi(\{x_i : i \in I\}) \nu^n(dx_1 \dots dx_n).$$

Grouping subsets by sizes yields

$$\begin{aligned} \mathbb{E} \sum_{Y \in \mathbf{X}} \phi(Y) &= e^{-\nu(D)} \sum_{n \geq 0} \frac{1}{n!} \sum_{\substack{m=0 \\ |I|=m}}^n \sum_{I \subset \{1, \dots, n\}} \nu(D)^{n-m} \\ &\quad \int_{D^m} \phi(\{x_1, \dots, x_m\}) \nu^m(dx_1 \dots dx_m). \end{aligned}$$

The summand does not depend on I . Therefore the sum over I is just equal to the number of subsets of m elements from a n -size set:

$$\begin{aligned}
& \mathbb{E} \sum_{Y \in \mathbf{X}} \phi(Y) \\
&= e^{-\nu(D)} \sum_{n \geq 0} \frac{1}{n!} \sum_{m=0}^n \frac{n! \nu(D)^{n-m}}{m!(n-m)!} \int_{D^m} \phi(\{x_1, \dots, x_m\}) \nu^m(dx_1 \dots dx_m), \\
&= e^{-\nu(D)} \sum_{m \geq 0} \frac{e^{\nu(D)}}{m!} \int_{D^m} \phi(\{x_1, \dots, x_m\}) \nu^m(dx_1 \dots dx_m), \\
&= e^{\nu(D)} \mathbb{E} \phi(\mathbf{X}).
\end{aligned}$$

□

C A random T -tessellation is determined by its Papan-gelou conditional intensities

Below a proof of Proposition 1 is given. From equations (7–8), it follows that

$$\begin{cases} h(sT) &= \lambda_s(s, T)h(T), \\ h(fT) &= \lambda_f(f, T)h(T), \end{cases}$$

for μ -almost all $T \in \mathcal{T}$ and almost all $s \in \mathbb{S}_T$ and all $f \in \mathbb{F}_T$. That is there exists a null set $E_s \subset \mathcal{C}_s$ with respect to the measure $d\mu(dT)$ such that $h(sT) = \lambda_s(s, T)h(T)$ for all $(s, T) \notin E_s$. Similarly there exists a null set $E_f \subset \mathcal{C}_f$ with respect to the measure $d\mu(dT)$ such that $h(fT) = \lambda_f(f, T)h(T)$ for all $(f, T) \notin E_f$.

Let

$$G = \left\{ T \in \mathcal{T} : \exists \tilde{T} \in \bigcup_{\tilde{L} \subset L(T)} \mathcal{T}(\tilde{L}), \exists o \in \mathbb{M}_{\tilde{T}} \cup \mathbb{F}_{\tilde{T}}, (o^{-1}, o\tilde{T}) \in E_s \cup E_f \right\},$$

where $L(T)$ is the pattern of lines supporting the segments of T . The union

$$\bigcup_{\tilde{L} \subset L(T)} \mathcal{T}(\tilde{L})$$

is the subset of T -tessellations that is spanned by applying successive merges and flips to T (remember that a merge remove a line supporting a segment and that a flip does not change the supporting lines).

Let $T \in \mathcal{T}$. There exists a sequence of splits and flips o_n, \dots, o_1 which transform gradually the empty tessellation D into T : $D = T_n, o_n T_n = T_{n-1}, \dots, o_2 T_2 = T_1, o_1 T_1 = T_0$ with $T_0 = T$. Note that for $i = 1, \dots, n$, $o_i \in \mathbb{S}_{T_i} \cup \mathbb{F}_{T_i}$ i.e. $o_i^{-1} \in \mathbb{M}_{T_{i-1}} \cup \mathbb{F}_{T_{i-1}}$. Suppose that $T \notin G$.

$$T_0 \notin G \Leftrightarrow \forall \tilde{T} \in \bigcup_{\tilde{L} \subset L(T)} \mathcal{T}(\tilde{L}), \forall o \in \mathbb{M}_{\tilde{T}} \cup \mathbb{F}_{\tilde{T}}, (o^{-1}, o\tilde{T}) \notin E_s \cup E_f.$$

The right-hand side statement holds for $\tilde{T} = T_0$. It follows

$$T_0 \notin G \Rightarrow \forall o \in \mathbb{M}_{T_0} \cup \mathbb{F}_{T_0}, (o^{-1}, oT_0) \notin E_s \cup E_f.$$

The right-hand side statement holds for $o = o_1$:

$$\begin{aligned} T_0 \notin G &\Rightarrow (o_1^{-1}, o_1 T_0) \notin E_s \cup E_f, \\ &\Rightarrow h(T_0) = \lambda_{t_1}(o_1, o_1^{-1} T_0) h(o_1^{-1} T_0), \\ &\Rightarrow h(T_0) = \lambda_{t_1}(o_1, T_1) h(T_1), \end{aligned}$$

where for $i = 1, \dots, n$

$$t_i = \begin{cases} s & \text{if } o_i \text{ is a split,} \\ f & \text{if } o_i \text{ is a flip.} \end{cases}$$

Repeating this computation, one gets

$$h(T) = h(D) \prod_{i=1}^n \lambda_{t_i}(o_i, T_i). \quad (46)$$

Now we show that G is a μ -null set. Let

$$\begin{aligned} G_s &= \left\{ T \in \mathcal{T} : \exists \tilde{T} \in \bigcup_{\tilde{L} \subset L(T)} \mathcal{T}(\tilde{L}), \exists o \in \mathbb{M}_{\tilde{T}}, (o^{-1}, o\tilde{T}) \in E_s \right\}, \\ G_f &= \left\{ T \in \mathcal{T} : \exists \tilde{T} \in \bigcup_{\tilde{L} \subset L(T)} \mathcal{T}(\tilde{L}), \exists o \in \mathbb{F}_{\tilde{T}}, (o^{-1}, o\tilde{T}) \in E_f \right\}. \end{aligned}$$

Obviously $G = G_s \cup G_f$. Let

$$F_s = \left\{ \tilde{T} \in \mathcal{T} : \exists m \in \mathbb{M}_{\tilde{T}}, (m^{-1}, m\tilde{T}) \in E_s \right\}.$$

We have

$$T \in G_s \Leftrightarrow \exists \tilde{T} \in \bigcup_{\tilde{L} \subset L(T)} \mathcal{T}(\tilde{L}), \tilde{T} \in F_s.$$

Therefore

$$\mu(G_s) \leq \mathbb{E} \sum_{\tilde{T} \in \bigcup_{\tilde{L} \subset L(\mathbf{T})} \mathcal{T}(\tilde{L})} I_{F_s}(\tilde{T}), \quad \mathbf{T} \sim \mu$$

Using Definition 1, one can rewrite the right-hand side as a mean with respect to the unit Poisson line process \mathbf{L} :

$$\begin{aligned} \mu(G_s) &\leq \text{cte } \mathbb{E} \sum_{T \in \mathcal{T}(\mathbf{L})} \sum_{\tilde{T} \in \bigcup_{\tilde{L} \subset L(T)} \mathcal{T}(\tilde{L})} I_{F_s}(\tilde{T}), \\ &\leq \text{cte } \mathbb{E} \sum_{T \in \mathcal{T}(\mathbf{L})} \sum_{\tilde{T} \in \bigcup_{\tilde{L} \subset L} \mathcal{T}(\tilde{L})} I_{F_s}(\tilde{T}), \\ &\leq \text{cte } \mathbb{E} |\mathcal{T}(\mathbf{L})| \sum_{\tilde{T} \in \bigcup_{\tilde{L} \subset L} \mathcal{T}(\tilde{L})} I_{F_s}(\tilde{T}). \end{aligned}$$

The constant before the mean is the unknown normalizing constant of the distribution μ . Now let us show that the random variable

$$\sum_{\tilde{T} \in \bigcup_{\tilde{L} \subset \mathbf{L}} \mathcal{T}(\tilde{L})} I_{F_s}(\tilde{T})$$

is null almost surely. Since it is non-negative, it is sufficient to show that its mean is null. Since the union is disjoint, we have

$$\mathbb{E} \sum_{\tilde{T} \in \bigcup_{\tilde{L} \subset \mathbf{L}} \mathcal{T}(\tilde{L})} I_{F_s}(\tilde{T}) = \mathbb{E} \sum_{\tilde{L} \subset \mathbf{L}} \sum_{\tilde{T} \in \mathcal{T}(\tilde{L})} I_{F_s}(\tilde{T}).$$

Applying Lemma 2 (see Appendix B), we can rewrite the mean of the sum on the line subsets:

$$\mathbb{E} \sum_{\tilde{T} \in \bigcup_{\tilde{L} \subset \mathbf{L}} \mathcal{T}(\tilde{L})} I_{F_s}(\tilde{T}) = \exp(b(D)/\pi) \mathbb{E} \sum_{\tilde{T} \in \mathcal{T}(\mathbf{L})} I_{F_s}(\tilde{T}),$$

where $b(D)$ is the perimeter of the domain D . Therefore it follows from Definition 1

$$\mathbb{E} \sum_{\tilde{T} \in \bigcup_{\tilde{L} \subset \mathbf{L}} \mathcal{T}(\tilde{L})} I_{F_s}(\tilde{T}) = \exp(b(D)/\pi) \mu(F_s)$$

From the definition of F_s we have

$$\begin{aligned} \mu(F_s) &\leq \mathbb{E} \sum_{m \in \mathbb{M}_{\mathbf{T}}} I_{E_s}(m^{-1}, mT), \quad \mathbf{T} \sim \mu, \\ &\leq \int_{\mathcal{T}} \int_{\mathbb{S}_T} I_{E_s}(s, T) ds \mu(dT), \\ &\leq 0. \end{aligned}$$

Hence

$$\mathbb{E} \sum_{\tilde{T} \in \bigcup_{\tilde{L} \subset \mathbf{L}} \mathcal{T}(\tilde{L})} I_{F_s}(\tilde{T}) = 0,$$

and

$$\sum_{\tilde{T} \in \bigcup_{\tilde{L} \subset \mathbf{L}} \mathcal{T}(\tilde{L})} I_{F_s}(\tilde{T})$$

is null almost surely. In turn, this implies that

$$\mathbb{E} |\mathcal{T}(\mathbf{L})| \sum_{\tilde{T} \in \bigcup_{\tilde{L} \subset \mathbf{L}} \mathcal{T}(\tilde{L})} I_{F_s}(\tilde{T}) = 0$$

and that $\mu(G_s) = 0$. The proof for $\mu(G_f) = 0$ is similar.

In Equation (46), the density h is determined only up to the multiplicative constant $h(D)$. It should be noticed that $h(D)$ is necessarily positive (otherwise h would be null on \mathcal{T} due to its hereditariness). The constant $h(D)$ is determined from the fact that the integral of h on \mathcal{T} is equal to 1.